

PATH-WISE ESTIMATORS AND CROSS-PATH REGRESSIONS: AN APPLICATION TO EVALUATING PORTFOLIO STRATEGIES

Martin B. Haugh

martin.b.haugh@gmail.com.

Ashish Jain

Columbia Business School, Uris Hall
3022 Broadway
New York, NY 10027, U.S.A.

ABSTRACT

Recently developed dual techniques allow us to evaluate a given sub-optimal dynamic portfolio policy by using the policy to construct an upper bound on the optimal value function. Moreover, when the policy is in fact optimal, the upper bound coincides with the optimal value function. Since it is easy to construct a lower bound by simulating the given policy, we may use the distance from the lower bound to the upper bound to assess the quality of the policy. One of the difficulties that arises when computing the upper bound, however, is that we need to know the sub-optimal policy's value function and its partial derivatives with respect to all state variables. If these quantities are not available analytically, then an alternative upper bound can still be computed but it is less satisfying from a theoretical perspective. In this paper we show how path-wise Monte-Carlo estimators together with the cross-path regression approach can be used to estimate the sub-optimal value function and its derivatives, thereby enabling us to compute the more theoretically satisfying upper bound on the optimal value function.

1 INTRODUCTION

Haugh, Kogan and Wang (2006, hereafter HKW) recently developed a dual-based technique for the evaluation of sub-optimal dynamic portfolio policies in a multi-dimensional diffusion setting in the presence of portfolio constraints and incomplete markets. The idea behind this technique, which is based on the earlier theoretical work of Xu (1990), Shreve and Xu (1992a, 1992b), Cvitanic and Karatzas (1992) and others, is to use the given sub-optimal policy to construct a *fictitious market* that is complete and unconstrained. By choosing the fictitious market in an appropriate manner, it can be shown that the optimal value function in the fictitious market provides an upper bound on the optimal value function in the original constrained and incomplete market. Since the standard martingale methods of Cox and

Huang (1989) and Karatzas, Lehocky and Shreve (1987) can then be used to compute the optimal value function in the fictitious market, HKW's approach allows us to go from a sub-optimal policy to an upper bound on the optimal value function.

Since it is easy to construct a lower bound by simply simulating the given policy, we may use the duality gap, i.e. the distance from the lower bound to the upper bound, to assess the quality of the portfolio policy. While HKW show that the duality gap will be zero if the sub-optimal policy is indeed optimal, the computation of the upper bound requires knowledge of the sub-optimal policy's value function and its derivatives with respect to any state variables in the system. For the numerical examples that HKW considered, this knowledge was unavailable and so they computed an alternative upper bound instead. This alternative upper bound proved to perform very well in that it was close to the optimal value function when the sub-optimal policy that generated it was also close to optimal. However, it suffered from some theoretical setbacks. In particular, even when the alternative upper bound was generated by the optimal policy itself, it was no longer guaranteed to coincide with the optimal value function.

Haugh and Jain (2007) studied the dual approach in further detail and, for certain classes of portfolio strategies, solved analytically for the corresponding value functions and their derivatives. They therefore succeeded in constructing the more theoretically satisfying upper bound originally proposed by HKW. However, when analytic solutions for the suboptimal value function are unavailable, the question remains as to whether or not the more theoretically satisfying upper bound can be computed. In this paper, we answer this question in the affirmative. In particular, we show how path-wise estimators and cross-path regressions may be successively used for estimating the sub-optimal policy's value function and its derivatives.

The remainder of this paper is organized as follows. In Section 2 we review the dual approach for evaluating sub-optimal strategies and explain why the suboptimal value

function and its derivatives are required for computing the more theoretically satisfying upper bound. In Section 3 we show how path-wise Monte-Carlo estimators and cross-path regressions can be used to estimate these quantities. Numerical results are presented in Section 4 and we conclude in Section 5.

2 THE DUAL APPROACH TO PORTFOLIO EVALUATION

In this section we briefly review HKW's dual-based technique for computing an upper bound on the optimal value function. Beginning with a sub-optimal strategy, we will see how this strategy can be used to define a fictitious market where the portfolio optimization problem is easily solved. Due to space constraints, we will not concern ourselves with too many technical details here and instead refer the reader to HKW (2007) for a more thorough description of the dual-based evaluation technique, and to Karatzas and Shreve (1997) for the theoretical foundations of the fictitious-market based duality approach to portfolio optimization.

We assume there are N risky assets and a single risk-free asset available in the economy. The time t vector of risky asset prices is denoted by $P_t = (P_t^{(1)}, \dots, P_t^{(N)})$ and the instantaneously risk-free rate of return is denoted by r_t . Security price dynamics are driven by the M -dimensional vector of state variables, X_t , so that

$$\begin{aligned} r_t &= r(X_t) \\ dP_t &= P_t[\mu_P(X_t)dt + \Sigma_P dB_t] \\ dX_t &= \mu_X(X_t)dt + \Sigma_X dB_t \end{aligned} \quad (1)$$

where $X_0 = 0$, $B_t = (B_{1t}, \dots, B_{Nt})$ is a vector of N independent Brownian motions, μ_P and μ_X are N and M dimensional drift vectors, and Σ_P and Σ_X are constant diffusion matrices of dimensions N by N and M by N , respectively. We can then define a process, η_t , as

$$\eta_t = \Sigma_{P_t}^{-1}(\mu_{P_t} - r_t).$$

In a market without portfolio constraints, η_t corresponds to the market price of risk process. See Duffie (1996) for further details.

A portfolio consists of positions in the N risky assets and the risk-free cash account. We denote the proportional holdings of the risky assets in the total portfolio value by $\theta_t = (\theta_{1t}, \dots, \theta_{Nt})$. The proportion in the risk-free asset is then given by $(1 - \theta_t^T \mathbf{1})$ where $\mathbf{1}$ is the unit vector of length N . The value of the portfolio, W_t , then has the following

dynamics

$$\frac{dW_t}{W_t} = \left[(1 - \theta_t^T \mathbf{1})r_t + \theta_t^T \mu_P(X_t) \right] dt + \theta_t^T \Sigma_P dB_t. \quad (2)$$

We assume that the proportional holdings in the portfolio are restricted to lie in a closed convex set, K , that contains the zero vector. In particular, we assume that

$$\theta_t \in K. \quad (3)$$

The investor's dynamic portfolio optimization problem is to solve for the value function V_0 at $t = 0$ where

$$V_0 \equiv \sup_{\theta_t} E_0[U(W_T)] \quad (4)$$

subject to constraints (1), (2) and (3).

The duality approach of HKW for analyzing the quality of a suboptimal strategy was to use the strategy to construct a lower and upper bound on the true value function. If the difference between the two bounds is large, i.e. the duality gap is wide, then it suggested that the suboptimal policy was not close to the optimal solution. If the duality gap is small, then (i) we know that the suboptimal strategy is close to optimal and (ii) we know approximately the optimal value function.

Starting with the portfolio optimization problem above, we can define a fictitious problem based on a different complete financial market and without the portfolio constraints. First we define the *support function* of \mathbf{K} , $\delta(\cdot) : \mathbb{R}^N \rightarrow \mathbb{N} \cup \infty$, by setting

$$\delta(v) = \sup_{x \in \mathbf{K}} (-v^T x). \quad (5)$$

The effective domain of the support function is given by

$$\tilde{\mathbf{K}} = \{v \in \mathbf{K} : \delta(v) < \infty\}.$$

Because the constraint set \mathbf{K} is convex and contains zero, the support function is continuous and bounded from below on its effective domain $\tilde{\mathbf{K}}$. We then define the set \mathbf{D} of adapted \mathbb{R}^N valued processes to be

$$\mathbf{D} = \left\{ v_t, 0 \leq t \leq T : v_t \in \tilde{\mathbf{K}}, E_0 \left[\int_0^T \delta(v_t) dt \right] + E_0 \left[\int_0^T \|v_t\|^2 dt \right] < \infty \right\}. \quad (6)$$

For each process v in \mathbf{D} , we define a fictitious market $M^{(v)}$. In this market, one can trade the N stocks and the risk-free cash account. The diffusion matrix of stock returns in $M^{(v)}$ is the same as in the original market. However, the risk-free rate and the vector of expected stock returns

are different. In particular, the risk-free rate process and the market price of risk in the fictitious market are defined respectively by

$$r_t^{(v)} = r_t + \delta(v_t) \quad (7a)$$

$$\eta_t^{(v)} = \eta_t + \Sigma_P^{-1} v_t \quad (7b)$$

where $\delta(v)$ is the support function defined in (5). We assume that $\eta_t^{(v)}$ is square-integrable. Following Cox and Huang (1989), the state-price density process $\pi_t^{(v)}$ in the fictitious market is given by

$$\pi_t^{(v)} = \exp\left(-\int_0^t r_s^{(v)} ds - \frac{1}{2} \int_0^t \eta_s^{(v)\top} \eta_s^{(v)} ds - \int_0^t \eta_s^{(v)\top} dB_s\right) \quad (8)$$

and the vector of expected returns is given by

$$\mu_{P_t}^{(v)} = r_t^{(v)} + \Sigma_P \eta_t^{(v)}.$$

The dynamic portfolio choice problem in the fictitious market without position constraints can be equivalently formulated in a static form:

$$\begin{aligned} V^{(v)} &\equiv \sup_{\{W_T\}} E_0[U(W_T)] \quad \text{subject to} \\ E_0\left[\pi_T^{(v)} W_T\right] &\leq W_0. \quad (\mathcal{P}^{(v)}) \end{aligned} \quad (9)$$

Due to its static nature, the problem $(\mathcal{P}^{(v)})$ is easy to solve. For example, when the utility function is of the CRRA type with relative risk aversion γ so that $U(W) = W^{1-\gamma}/(1-\gamma)$, the corresponding value function in the fictitious market is given explicitly by

$$V_0^{(v)} = \frac{W_0^{1-\gamma}}{1-\gamma} E_0\left[\pi_T^{(v)\frac{\gamma-1}{\gamma}}\right]^\gamma. \quad (10)$$

Moreover, it is straightforward to show that the solution to $(\mathcal{P}^{(v)})$ provides an upper bound on the optimal value function in the original portfolio optimization problem. While one can pick any fictitious market from the admissible set \mathbf{D} to compute an upper bound, HKW showed how a given suboptimal strategy, $\tilde{\theta}_t$, may be used to select a particular $\hat{v}_t \in \mathbf{D}$. If the suboptimal strategy is in fact optimal, then the lower bound associated with the suboptimal strategy will equal the associated upper bound, thereby demonstrating its optimality.

Given an approximation to the optimal portfolio policy $\tilde{\theta}_t$, one can compute the corresponding approximation to the value function, \tilde{V}_t , defined as the conditional expectation of the utility of terminal wealth, under the portfolio policy $\tilde{\theta}_t$.

HKW then define $\tilde{\eta}_t$ as

$$\tilde{\eta}_t := -W_t \left(\frac{\partial_{WW} \tilde{V}_t}{\partial_W \tilde{V}_t} \right) \Sigma_P^\top \tilde{\theta}_t - \left(\partial_W \tilde{V}_t \right)^{-1} \Sigma_X^\top \left(\partial_{WX} \tilde{V}_t \right) \quad (11)$$

where ∂_W denotes the partial derivative with respect to W , and ∂_{WX} and ∂_{WW} are corresponding second partial derivatives. We then define \tilde{v}_t as a solution to (7b). This definition of $\tilde{\eta}_t$ (see HKW for the motivation behind this definition) links the sub-optimal strategy, $\tilde{\theta}_t$, to an upper bound on the optimal value function for the original portfolio optimization problem.

In the special but important case of a CRRA utility function the expression for $\tilde{\eta}_t$ simplifies. In particular, for a given trading strategy, $\tilde{\theta}_t$, and CRRA utility, the corresponding value function is of the following form

$$\tilde{V}_t = g(t, X_t) \frac{W_t^{1-\gamma}}{1-\gamma} \quad (12)$$

for some function, $g(\cdot, \cdot)$.

Hence, the market price of risk process in the dual problem simplifies to

$$\begin{aligned} \tilde{\eta}_t &= \gamma \Sigma_P^\top \tilde{\theta}_t - \frac{\Sigma_t^\top}{\tilde{V}_t} \left(\frac{\partial \tilde{V}_t}{\partial X_t} \right) \\ &= \gamma \Sigma_P^\top \tilde{\theta}_t - \frac{\Sigma_X^\top}{g(t, X_t)} \left(\frac{\partial g(t, X_t)}{\partial X_t} \right) \end{aligned} \quad (13)$$

where γ is the coefficient of relative risk aversion. Note that one only needs to compute the first derivative of the value function with respect to the state variables, X_t , to evaluate the second term in (13). This simplifies numerical implementation, since it is easier to estimate first-order than second-order partial derivatives of the value function. For general strategies we don't have an analytical solution for the value function and its derivatives. In the remainder of this paper we will use the phrases θ -Term and g -Term when referring to the first and second terms, respectively, on the right-hand-side of (13).

Obviously, $\tilde{\eta}_t$ is a candidate for the market price of risk in the fictitious market. However, there is no guarantee that $\tilde{\eta}_t$ and the corresponding process, \tilde{v}_t , belong to the feasible set \mathbf{D} defined by (6). In fact, for many important classes of problems the support function $\delta(v_t)$ may be infinite for some values of its argument. We therefore look for a price-of-risk process $\hat{\eta}_t \in \mathbf{D}$ that is "close" to $\tilde{\eta}_t$ by formulating a simple quadratic optimization problem. Depending on the portfolio constraints, this problem may be solved analytically. Otherwise, we solve it numerically at each discretization point on each simulated path of the underlying stochastic differential equations (SDE's). The lower bound is then computed by simulating the given

portfolio strategy. The same simulated paths of the SDE's are then used to estimate the upper bound given by (10). At each discretization point on each simulated path we solve a quadratic optimization problem to find the appropriate $\hat{\eta}_t \in \mathbf{D}$. See HKW for further details. We use a simple Euler scheme when simulating the underlying SDE's. See Kloeden and Platen (1992) for more sophisticated techniques.

In their numerical experiments, HKW did not have an analytic expression available for the g -Term so they were unable to compute the theoretical upper bound that they proposed. Instead, they used only the θ -Term to compute their upper bounds. While Haugh and Jain (2007) were able to compute both terms analytically for certain classes of portfolio strategies, this is not possible in general. Approximating the g -Term is therefore of interest and is the focus of this paper. In particular, we will show how path-wise Monte Carlo estimators and *cross-path regressions* can be used to estimate the g -Term.

3 PATH-WISE ESTIMATORS AND CROSS-PATH REGRESSIONS

We assume now a particular form for the dynamics that our security prices satisfy. These are the same dynamics assumed by HKW and Haugh and Jain (2007). We have

$$\begin{aligned} r_t &= r \\ dP_t &= P_t[(\mu_0 + X_t\mu_1)dt + \Sigma_P dB_t^P] \\ dX_t &= -kX_t dt + \Sigma_X dB_t^P. \end{aligned} \quad (14)$$

The first equation gives the risk free rate which is assumed constant in our numerical results. The second equation specifies the dynamics of the three traded risky securities. The diffusion matrix Σ_X is of size 1 by 4 and coincides with the last row of matrix Σ_P . The vectors μ_0 and μ_1 define the drift vector for the risky securities. The third equation specifies the dynamics of the state variable, X_t , whose initial value is set to zero in all of the numerical examples.

We now derive the path-wise estimator for the sub-optimal value function and its derivatives. (Glasserman (2003) is a good source and contains many financial applications of path-wise estimators.) We then briefly describe the cross-path regression approach for evaluating this estimator.

3.1 The Path-Wise Estimator

Let $\theta_t = \theta(X_t)$ denote the given sub-optimal policy that depends explicitly on the state variable(s), X_t . We wish to estimate its corresponding value function, V_t , and its first derivative, $\frac{\partial V_t}{\partial X_t}$. The main idea behind the path-wise estimator is the interchange of the differentiation and expectation operators so that

$$\frac{d}{dX_t} E_t[h(X_t)] = E_t \left[\frac{d}{dX_t} h(X_t) \right] \quad (15)$$

where $h(\cdot)$ is a generic function. The expression $\frac{d}{dX_t} h(X_t)$ is the path-wise derivative of $h(X_t)$ with respect to the state variable, X_t . While some technical conditions are required in order to justify the above interchange of differentiation and expectation, they do not present a problem in our application in this paper.

In the case of CRRA utility, we can write

$$V_t = E_t^\theta \left[U(W_T) \right] = E_t^\theta \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right] \quad (16)$$

where we use the $E_t^\theta[\cdot]$ notation to acknowledge the dependence of the terminal wealth, W_T , on the strategy, θ . Interchanging the expectation and differential operator in order to compute the first derivative, we obtain

$$\frac{\partial V_t}{\partial X_t} = \frac{1}{1-\gamma} E_t^\theta \left[\frac{\partial W_T^{1-\gamma}}{\partial X_t} \right]. \quad (17)$$

Applying Itô's Lemma to $\log(W_t)$ and using (2) we obtain

$$\begin{aligned} W_T &= W_t \exp \left(\int_t^T \left((1 - \theta_s^\top \mathbf{1})r + \theta_s^\top (\mu_0 + \mu_1 X_s) \right. \right. \\ &\quad \left. \left. - \frac{1}{2} (\theta_s^\top \Sigma_P \Sigma_P^\top \theta_s) \right) ds + \int_t^T \theta_s^\top \Sigma_P dB_s^P \right) \end{aligned} \quad (18)$$

so that

$$\begin{aligned} \frac{\partial W_T^{1-\gamma}}{\partial X_t} &= (1-\gamma) W_T^{1-\gamma} \frac{\partial}{\partial X_t} \left(\int_t^T \left((1 - \theta_s^\top \mathbf{1})r \right. \right. \\ &\quad \left. \left. + \theta_s^\top (\mu_0 + \mu_1 X_s) - \frac{1}{2} \theta_s^\top \Sigma_P \Sigma_P^\top \theta_s \right) ds + \int_t^T \theta_s^\top \Sigma_P dB_s^P \right) \end{aligned}$$

We can therefore express the value function's first derivative with respect to X_t as

$$\begin{aligned} \frac{\partial V_t}{\partial X_t} &= E_t^\theta \left[W_T^{1-\gamma} \frac{\partial}{\partial X_t} \left(\int_t^T \left((1 - \theta_s^\top \mathbf{1})r + \theta_s^\top (\mu_0 + \mu_1 X_s) \right) ds \right. \right. \\ &\quad \left. \left. - \frac{1}{2} \int_t^T (\theta_s^\top \Sigma_P \Sigma_P^\top \theta_s) ds + \int_t^T \theta_s^\top \Sigma_P dB_s^P \right) \right] \end{aligned}$$

In our numerical experiments we simulate the asset return dynamics over a discrete time grid, i.e. we use a standard Euler scheme to simulate the SDE's. We let $0 = t_0 < t_1 < \dots < t_m = T$ be a partition of the time interval $[0, T]$ into m equal segments of length Δt so that $t_i = iT/m$ for $i = 0, 1, \dots, m$. If the current value of t is $t = t_i$, then the

discretized integral inside the differential operator can be expressed as

$$\begin{aligned}
 & \int_{t_i}^T ((1 - \theta_s^\top \mathbf{1})r + \theta_s^\top (\mu_0 + \mu_1 X_s)) ds \\
 & - \frac{1}{2} \int_{t_i}^T (\theta_s^\top \Sigma_P \Sigma_P^\top \theta_s) ds + \int_{t_i}^T \theta_s^\top \Sigma_P dB_s^P \\
 = & \sum_{j=i}^{m-1} ((1 - \theta_j^\top \mathbf{1})r + \theta_j^\top (\mu_0 + \mu_1 X_j)) \Delta t \quad (19) \\
 & - \frac{1}{2} \sum_{j=i}^{m-1} (\theta_j^\top \Sigma_P \Sigma_P^\top \theta_j) \Delta t + \sum_{j=i}^{m-1} \theta_j^\top \Sigma_P Z_{j+1} \sqrt{\Delta t}
 \end{aligned}$$

where Z_j 's form a sequence of IID standard normal random variables. We can easily solve the SDE describing the dynamics of X_t to obtain

$$X_t = \exp(-k(t-s)) \left(X_s + \int_s^t \exp(k(u-s)) \Sigma_X dB_u^P \right).$$

For $j > i$, we therefore obtain

$$\frac{\partial X_{t_j}}{\partial X_{t_i}} = \exp(-k(t_j - t_i)) = \exp(-k(j-i)\Delta t) \quad (20)$$

and, by the chain rule,

$$\frac{\partial \theta_j}{\partial X_i} = \frac{\partial \theta_j}{\partial X_j} \frac{\partial X_j}{\partial X_i}. \quad (21)$$

Applying (20) and (21) to differentiate (19) yields

$$\begin{aligned}
 & \frac{\partial}{\partial X_{t_i}} \left(\sum_{j=i}^{m-1} ((1 - \theta_j^\top \mathbf{1})r + \theta_j^\top (\mu_0 + \mu_1 X_j)) \Delta t \right. \\
 & \left. - \frac{1}{2} \sum_{j=i}^{m-1} (\theta_j^\top \Sigma_P \Sigma_P^\top \theta_j) \Delta t + \sum_{j=i}^{m-1} \theta_j^\top \Sigma_P Z_{j+1} \sqrt{\Delta t} \right) \\
 = & \sum_{j=i}^{m-1} \left(\left(-\frac{\partial \theta_j^\top}{\partial X_j} \mathbf{1} \exp(-k(j-i)\Delta t) \right) r \right. \\
 & \left. + \frac{\partial \theta_j^\top}{\partial X_j} \exp(-k(j-i)\Delta t) (\mu_0 + \mu_1 X_j) \right) \Delta t \\
 & + \left(\sum_{j=i}^{m-1} \theta_j^\top \mu_1 \exp(-k(j-i)\Delta t) \right) \Delta t \\
 & - \frac{1}{2} \sum_{j=i}^{m-1} \left(2\theta_j^\top \exp(-k(j-i)\Delta t) \Sigma_P \Sigma_P^\top \frac{\partial \theta_j}{\partial X_j} \right) \Delta t
 \end{aligned}$$

$$+ \sum_{j=i}^{m-1} \left(\frac{\partial \theta_j^\top}{\partial X_j} \exp(-k(j-i)\Delta t) \Sigma_P Z_{j+1} \sqrt{\Delta t} \right) =: PW. \quad (22)$$

Hence the first derivative of the policy's value function can be expressed as

$$\frac{\partial V_{t_i}}{\partial X_{t_i}} = E_{t_i}^\theta \left[W_T^{1-\gamma} PW \right] \quad (23)$$

To compute the g -Term in (13) we need the ratio of $g(t, X_t)$ and its first derivative with respect to the state variable, X_t . This ratio is estimated by estimating the numerator and denominator separately. Equations (12), (16) and (23) imply we can estimate these quantities using

$$g(t_i, X_{t_i}) = E_{t_i}^\theta \left[\frac{W_T^{1-\gamma}}{W_{t_i}^{1-\gamma}} \right] \quad (24)$$

$$\frac{\partial g(t_i, X_{t_i})}{\partial X_{t_i}} = E_{t_i}^\theta \left[\frac{W_T^{1-\gamma}}{W_{t_i}^{1-\gamma}} PW (1-\gamma) \right]. \quad (25)$$

Because we need to estimate these quantities at all discretization points, t_i , and at all realizations of the state variable, X_{t_i} , we need an efficient method for computing these estimates. This leads to the use of the cross-path regression approach which we now briefly discuss.

3.2 The Cross-Path Regression Approach

The cross-path regression technique was introduced by Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001) in the context of pricing high-dimensional American options. Their approach was to simulate a large number of sample paths of the underlying state variables and use the standard dynamic programming approach of value iteration to estimate the option value at each time point on each simulated path. Rather than directly estimating the option value, they estimated the *continuation value*, i.e. the value of the option conditional on it not being exercised at the current time period. The key technique in the cross-path regression approach was to estimate all of the continuation values at a given time period through a *single* linear regression that utilized the cross-sectional information in the simulated paths.

Brandt et al (2005) extended this idea to dynamic portfolio optimization problems with incomplete markets. Haugh, Kogan and Wu (2006) considered the same problem but with the added complication of imposing no-borrowing constraints. In this paper we will use the same cross-path regression approach to estimate each of the quantities in (24) and (25). Due to space

constraints we will only outline the details of our algorithm:

The Cross-Path Regression Algorithm

1. Simulate N paths of the state variable, X_t , and the security prices, P_t , at the time points $0 < t_1 < \dots < t_m = T$.
2. For $i = 1, \dots, m$, compute the trading strategy, θ_{t_i} , and the corresponding wealth, W_{t_i} , on each path at time t_i . The corresponding value function at time T satisfies $V_T = W_T^{1-\gamma}/(1-\gamma)$.
3. For $i = (m-1), \dots, 1$, estimate $E_{t_i}^\theta [(W_T/W_{t_i})^{1-\gamma}]$ and $E_{t_i}^\theta [(W_T/W_{t_i})^{1-\gamma}PW(1-\gamma)]$ on each path by regressing $(W_T/W_{t_i})^{1-\gamma}$ and $(W_T/W_{t_i})^{1-\gamma}PW(1-\gamma)$, respectively, on X_{t_i} and some higher powers of X_{t_i} .

The estimated regressions at time t_i can now be used to estimate the ratio $\frac{\partial g_{t_i}}{\partial X_{t_i}}/g_{t_i}$ that we need for the g -Term on the right-hand side of (13). Note also that because linear regressions can be performed very quickly, the computational requirements for the cross-path regression algorithm are not particularly demanding.

4 NUMERICAL RESULTS

We now apply the cross-path simulation algorithm to three different portfolio strategies, namely the static, generalized buy-and-hold (GBH) and myopic portfolio strategies. Security price and state dynamics follow (14) with parameters taken from HKW. In particular, the parameter set we use in this paper (see Table 1) coincides with Parameter Set #2 in HKW which was taken, in turn, from Lynch (2001). This parametrization assumes the three risky securities are in fact size-sorted portfolios with the term spread (the difference in yields between 20-year and 1-month Treasury securities) as the state variable. All our results assume a horizon of $T = 5$ years, and an Euler scheme with 100 time periods per year was used for simulating the underlying SDE's. At each time step t , we used $\{1, X_t, \dots, X_t^m\}$ as the dependent variables in each of the two linear regressions required to estimate the numerator and denominator of the g -Term. We generally took m equal to 5 or 6.

4.1 The Static and GBH Trading Strategies

We first consider the static and GBH strategies because analytic solutions for the g -Term are available (see Haugh and Jain 2007) in those cases. This means we can compute the estimation error directly for these strategies.

k	μ_0	μ_1	Σ_P			
1.671	0.081	0.046	0.186	0.000	0.000	0.000
	0.110	0.070	0.227	0.082	0.000	0.000
	0.130	0.086	0.251	0.139	0.069	0.000
	0.000	0.000	-0.017	0.149	0.058	1.725

Table 1: **Calibrated Model Parameters.**

The static trading strategy is the well-known *constant proportion* trading strategy that, at time t , invests a fixed proportion of the available wealth, W_t , in each of the risky securities. It satisfies

$$\theta_t^{static} = \arg \max_{\theta} (\mu_0^T - r)\theta - \frac{1}{2}\gamma\theta^T \Sigma_P \Sigma_P^T \theta, \quad (26)$$

so $\theta_t^{static} = \theta^{st}$ is a constant vector that does not vary with time. It solves the investor's dynamic portfolio optimization problem when the instantaneous first and second moments of security returns are assumed to equal to their long-term averages.

The GBH trading strategy was introduced in Haugh and Jain (2007). It is defined as the optimal trading strategy where the terminal wealth depends only on the terminal security prices. In particular, it is the optimal trading strategy resulting in a non-path-dependent wealth. It may be shown (see Haugh and Jain 2007) that the trading strategy that replicates the optimal GBH terminal wealth is a function of time only, and that it does not depend on the path of the state variable, X_t .

Because the static and GBH trading strategies are state-variable independent we have

$$\frac{\partial \theta_j}{\partial X_j} = 0$$

at all times, t . The path-wise estimator of (22) therefore reduces to

$$PW = \sum_{j=i}^{m-1} \left(\theta_j^\top \mu_1 \exp(-k(j-i)\Delta t) \right) \Delta t$$

where $\theta_j^\top = \theta^{st}$ for the static strategy, and is time dependent but state-variable independent, for the GBH strategy.

In order to study the performance of the cross-path regression approach we compute two types of errors for

each time, t_i :

$$\text{Error}_1 := \frac{1}{N} \sum_{j=1}^N \frac{|\hat{f}(X_{t_i}^j) - f(X_{t_i}^j)|}{f(X_{t_i}^j)} \quad (27)$$

$$\text{Error}_2 := \sqrt{\frac{\sum_{j=1}^N \left(1 - \hat{f}(X_{t_i}^j)/f(X_{t_i}^j)\right)^2}{N}} \quad (28)$$

where N is the number of sample paths, $f(\cdot)$ represents the function we wish to estimate and $\hat{f}(\cdot)$ is the cross-path regression estimator of function $f(\cdot)$. Table 2 displays Error₁ and Error₂ for the θ -Term and g -Term. Recall that Haugh and Jain (2007) obtained analytic solutions for the g -Term in the case of the static and GBH strategies. We assume the coefficient of relative risk aversion satisfies $\gamma = 1.5$ and that $N = 1$ million paths X_t and P_t were simulated. The errors were estimated at the mid-point of the time horizon, i.e., where $t_i = 2.5$ years so that $i = 250$.

	Error ₁	Error ₂
$g(X_{t_i})$	0.0011	0.0018
$\frac{\partial g(X_{t_i})}{\partial X_{t_i}}$	0.0013	0.0017
$\frac{\partial g(X_{t_i})}{\partial X_{t_i}} / g(X_{t_i})$	0.0014	0.0020

Table 2: Cross-Path Regression Errors

We see that in each case the estimated errors are less than 0.2% suggesting that the regression-based estimators are very accurate. We now examine the quality of the upper bound when we use these regression-based estimates of the g -Term.

Table 3 displays the upper bound on the optimal value function that is generated using the static and GBH strategies, respectively. In each case we estimate the upper bound using both the regression-based estimate as well as the analytic solution for the g -Term. We also include the upper bound for the static strategy that was estimated by HKW using only the θ -Term from (13), as well as the true optimal value function, which can be computed in this case. Approximate 95% confidence intervals are in parentheses.

We see that the upper bounds obtained using the regression approach are almost identical to those obtained from the analytical solution and so we can conclude, at least for this model and the static and GBH strategies, that the cross-path regression algorithm works very well. While using the g -Term to compute the upper bound does not produce a significant improvement to the upper bound in absolute terms, it does produce a significant *relative* improvement when compared to the optimal value function.

	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$
UB_{Static}^{HKW}	9.20 (9.15, 9.24)	5.09 (5.04, 5.14)	3.47 (3.42, 3.53)
$UB_{Static}^{Analytic}$	9.14 (9.10, 9.19)	5.00 (4.95, 5.05)	3.38 (3.33 , 3.43)
$UB_{Static}^{Regression}$	9.14 (9.10, 9.19)	5.00 (4.95, 5.05)	3.38 (3.33 , 3.43)
$UB_{GBH}^{Analytic}$	9.19 (9.14, 9.24)	5.02 (4.94 , 5.10)	3.38 (3.30 , 3.47)
$UB_{GBH}^{Regression}$	9.19 (9.14, 9.24)	5.02 (4.94 , 5.10)	3.38 (3.30 , 3.47)
$V_{Optimal}$	9.02	4.92	3.33

Table 3: Regression-Based Upper Bounds: Static and GBH Strategies

4.2 The Myopic Trading Strategy

We now consider the myopic trading strategy which is defined in the same way as the static policy, except the instantaneous moments of asset returns are fixed at their current values, as opposed to their long-run average values. In particular, the optimizing agent solves

$$\theta_t^{myopic} = \arg \max_{\theta} (\mu_{P_t}^T - r)\theta - \frac{1}{2}\gamma\theta^T \Sigma_P \Sigma_P^T \theta, \quad (29)$$

subject to $\theta \in K$.

In the absence of portfolio constraints, the myopic portfolio strategy satisfies

$$\theta_t^{myopic} = \frac{(\Sigma_P \Sigma_P^T)^{-1}(\mu_{P_t}^T - r)}{\gamma} \quad (30)$$

and so its derivative with respect to the state variable, X_t , is given by

$$\frac{\partial \theta_j^{myopic}}{\partial X_j} = \frac{(\Sigma_P \Sigma_P^T)^{-1} \mu_1}{\gamma}. \quad (31)$$

The path-wise estimator can therefore be computed explicitly using (22). Because an analytic expression for the g -Term is unavailable in this case, we assess the performance of the cross-path regression estimator by comparing the corresponding upper bound with the upper bound obtained by HKW who only used the θ -Term.

The results are displayed in Table 4 where again we report the optimal value function for comparison purposes. We see once again that using the regression based estimator of the g -Term produces a significant relative improvement to the quality of the upper bound. While not the focus of this paper, it is worth noting that although all three sub-

optimal strategies produce comparable upper bounds their corresponding lower bounds can vary significantly. This is not always the case. See HKW and Haugh and Jain (2007) for further examples.

We also mention at this point that it is no accident that the myopic and static upper bounds from HKW coincide. This occurs only in the case of unconstrained markets and when only the θ -Term is used to construct the upper bound. See HKW for further details.

	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$
UB_{Myopic}^{HKW}	9.20 (9.15, 9.24)	5.09 (5.04, 5.14)	3.47 (3.42, 3.53)
$UB_{Myopic}^{Regression}$	9.13 (9.09, 9.17)	4.99 (4.94, 5.04)	3.37 (3.32, 3.42)
$V_{Optimal}$	9.02	4.92	3.33

Table 4: **Regression-Based Upper Bounds: Myopic Strategy**

5 CONCLUSIONS AND FURTHER RESEARCH

In this paper we have shown how path-wise Monte-Carlo estimators and the cross-path regression approach could be used for estimating a given sub-optimal policy's value function and its derivatives. In particular, we have shown how the second term on the right-hand-side of (13) might be successfully estimated. This enabled us to compute the upper bound on the optimal value function that was originally proposed by HKW. We saw that our estimation algorithm was able to produce upper bounds that were comparable to the upper bounds that utilized the analytic expression, when it was available, for the second term. In the case of the myopic strategy, we succeeded in improving HKW's upper bound that utilized only the θ -Term. While the improvement was quite small in magnitude, it was still significant given how close the HKW bound was to the optimal solution. In conclusion, our results suggest that it should often be possible to estimate the theoretical upper bound originally proposed by HKW for generic sub-optimal strategies, even when the corresponding second term on the right-hand-side of (13) is not available analytically.

There are several possible directions for future research. First, we would like to extend the work to the case where portfolio constraints such as no-borrowing or no short-sales are imposed. These constraints are considered in HKW and in Haugh and Jain (2007). The simulation algorithm of this paper should easily carry over to constrained problems and so this extension should be very straightforward.

Second, we would like to apply this algorithm to other classes of sub-optimal strategies such as the strategies resulting from approximate dynamic programming (ADP) tech-

niques. (See for example, Brandt et al 2005 and Haugh, Kogan and Wu 2006.)

Finally, it would be of interest to find a problem where the upper bound generated by the first term only is significantly inferior to the upper bound generated by the first and second terms taken together. While we do not know of such a problem, it would serve to highlight the significance of being able to accurately estimate the second term.

REFERENCES

- Brandt, M.W., A. Goyal, P. Santa-Clara, and J.R. Stroud. 2005. A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability. *Review of Financial Studies* 18: 831-873.
- Cox, J., and C.-F. Huang. 1989. Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process. *Journal of Economic Theory* 49:33-83.
- Cvitanić, J., and I. Karatzas. 1992. Convex Duality in Constrained Portfolio Optimization. *Annals of Applied Probability* 2:767-818.
- Duffie, D. 1996. *Dynamic Asset Pricing Theory*. Princeton, New Jersey: Princeton University Press.
- Glasserman, P. 2003. *Monte Carlo Methods in Financial Engineering*. Springer, New York.
- Haugh, M.B. and A. Jain. 2007. On the Dual Approach to Portfolio Evaluation. Working paper.
- Haugh, M.B., L. Kogan, and J. Wang. 2003. Evaluating Portfolio Strategies: A Duality Approach. *Operations Research*. Vol.54.3. 4015-418.
- Haugh, M.B., L. Kogan and Z. Wu. 2006. Portfolio Optimization with Position Constraints: an Approximate Dynamic Programming Approach. Working paper. Columbia University and MIT.
- Karatzas, I., J.P. Lehocky and S.E. Shreve. 1987. Optimal Portfolio and Consumption Decisions for a "Small Investor" on a Finite Horizon. *SIAM J. Control Optimization* 25:1557-1586.
- Karatzas, I., and S.E. Shreve. 1997. *Methods of Mathematical Finance*. New York: Springer-Verlag.
- Kloeden, P., and E. Platen. 1992. *Numerical Solution of Stochastic Differential Equations*. Berlin: Springer-Verlag.
- Longstaff, F., and E. Schwartz. 2001. Valuing American Options by Simulation: A Simple Least-Squares Approach. *Review of Financial Studies* 14:113-147.
- Lynch, A. 2001. Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands Induced by Return Predictability. *Journal of Financial Economics*. 62 67-130.
- Shreve, S.E., and G.L. Xu. 1992a. A Duality Method for Optimal Consumption and Investment under Short-Selling Prohibition, Part I: General Market Coefficients. *Annals of Applied Probability* 2:8-112.

- Shreve, S.E., and G.L. Xu. 1992b. A Duality Method for Optimal Consumption and Investment under Short-Selling Prohibition, Part I: Constant Market Coefficients. *Annals of Applied Probability* 2:314-328.
- Tsitsiklis, J., and B. Van Roy. 2001. Regression Methods for Pricing Complex American-Style Options. *IEEE Transactions on Neural Networks* 12(4):694-703.
- Xu, G.L. 1990. A Duality Method for Optimal Consumption and Investment under Short-Selling Prohibition. Ph.D. Dissertation, Carnegie Mellon University.

AUTHOR BIOGRAPHY

Martin B. Haugh was an Assistant Professor in the Department of Industrial Engineering and Operations Research at Columbia University from 2002 to 2005. He currently works in the financial services industry. He may be contacted at <martin.b.haugh@gmail.com>.

Ashish Jain works in the financial services industry. The research for this paper was completed while he was a doctoral student at Columbia Business School. He may be contacted at <aj2109@columbia.edu>.